

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

**TRANSMITTAL
FORM**

(to be used for all correspondence after initial filing)

Total Number of Pages in This Submission

33

Application Number 10/580,792

Filing Date 28 April 2006

First Named Inventor Benjamin Liu

Art Unit

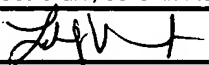
Examiner Name

Attorney Docket Number 42P22229

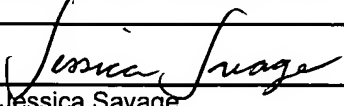
ENCLOSURES (Check all that apply)

<input type="checkbox"/> Fee Transmittal Form	<input type="checkbox"/> Drawing(s)	<input type="checkbox"/> After Allowance Communication to TC
<input type="checkbox"/> Fee Attached	<input type="checkbox"/> Licensing-related Papers	<input type="checkbox"/> Appeal Communication to Board of Appeals and Interferences
<input type="checkbox"/> Amendment/Reply	<input type="checkbox"/> Petition	<input type="checkbox"/> Appeal Communication to TC (Appeal Notice, Brief, Reply Brief)
<input type="checkbox"/> After Final	<input type="checkbox"/> Petition to Convert to a Provisional Application	<input type="checkbox"/> Proprietary Information
<input type="checkbox"/> Affidavits/declaration(s)	<input type="checkbox"/> Power of Attorney, Revocation	<input type="checkbox"/> Status Letter
<input type="checkbox"/> Extension of Time Request	<input type="checkbox"/> Change of Correspondence Address	<input checked="" type="checkbox"/> Other Enclosure(s) (please identify below): Return Receipt Postcard
<input type="checkbox"/> Express Abandonment Request	<input type="checkbox"/> Terminal Disclaimer	
<input type="checkbox"/> Information Disclosure Statement	<input type="checkbox"/> Request for Refund	
<input checked="" type="checkbox"/> Certified Copy of Priority Document(s)	<input type="checkbox"/> CD, Number of CD(s) _____	
<input type="checkbox"/> Reply to Missing Parts/Incomplete Application	<input type="checkbox"/> Landscape Table on CD	
<input type="checkbox"/> Reply to Missing Parts under 37 CFR 1.52 or 1.53	Remarks	Express Mail No. EV 841 071 555 US

SIGNATURE OF APPLICANT, ATTORNEY, OR AGENT

Firm Name	BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN, LLP		
	12400 Wilshire Boulevard, Seventh Floor, Los Angeles, California 90025-1030		
Signature			
Printed name	Lester J. Vincent		
Date	August 29, 2006	Reg. No.	31, 460

CERTIFICATE OF MAILING

I hereby certify that this correspondence is being deposited with the United States Postal Service with sufficient postage as first class mail in an envelope addressed to: M/S PCT, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450 on the date shown below:			
Signature			
Typed or printed name	Jessica Savage	Date	August 29, 2006

This collection of information is required by 37 CFR 1.5. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to 2 hours to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: M/S PCT, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

中华人民共和国国家知识产权局
STATE INTELLECTUAL PROPERTY OFFICE
OF THE PEOPLE'S REPUBLIC OF CHINA



证 明 CERTIFICATE

本证明之附件是向中国专利局作为受理局提交的下列国际申请副本
IS TO CERTIFY THAT ANNEXED HERETO IS A TRUE COPY OF THE BELOW
IDENTIFIED INTERNATIONAL APPLICATION THAT WAS FILED WITH THE
CHINESE PATENT OFFICE AS RECEIVING OFFICE

申请号:

PCT/CN2005/002305

INTERNATIONAL APPLICATION NUMBER

申请日:

23. DEC 2005 (23.12.2005)

INTERNATIONAL FILING DATE

名称:

CHANGING A SCHEDULER IN A VIRTUAL MACHINE MONITOR

INVENTION

CERTIFIED COPY OF
PRIORITY DOCUMENT

中华人民共和国国家知识产权局局长
COMMISSIONER OF THE STATE INTELLECTUAL PROPERTY
OFFICE OF THE PEOPLE'S REPUBLIC OF CHINA

二零零六年五月十七日

MAY 17 2006

HOLE COPY

PCT

REQUEST

The undersigned requests that the present international application be processed according to the Patent Cooperation Treaty.

For receiving Office use only

PCT/CN 2005 / 002305	
International Application No.	
23 DEC. 2005 (23.12.2005)	
International Filing Date	
RO/CN 中华人民共和国国家知识产权局 PCT International Application	
Name of receiving Office and "PCT International Application"	
Applicant's or agent's file reference (if desired) (12 characters maximum) FPEL05150069	

Box No. I TITLE OF INVENTION CHANGING A SCHEDULER IN A VIRTUAL MACHINE MONITOR	
Box No. II APPLICANT <input type="checkbox"/> This person is also inventor	
Name and address: (Family name followed by given name; for a legal entity, full official designation. The address must include postal code and name of country. The country of the address indicated in this Box is the applicant's State (that is, country) of residence if no State of residence is indicated below.) INTEL CORPORATION 2200 Mission College Blvd. Santa Clara, California 95052 United States of America	Telephone No. Facsimile No. Teleprinter No. Applicant's registration No. with the Office
State (that is, country) of nationality: US	State (that is, country) of residence: US
This person is applicant for the purposes of: <input type="checkbox"/> all designated States <input checked="" type="checkbox"/> all designated States except the United States of America <input type="checkbox"/> the United States of America only <input type="checkbox"/> the States indicated in the Supplemental Box	
Box No. III FURTHER APPLICANT(S) AND/OR (FURTHER) INVENTOR(S)	
Name and address: (Family name followed by given name; for a legal entity, full official designation. The address must include postal code and name of country. The country of the address indicated in this Box is the applicant's State (that is, country) of residence if no State of residence is indicated below.) LIU, Benjamin Room 302, No. 264 South Shanxi Road Shanghai 200001 P. R. of China	This person is: <input type="checkbox"/> applicant only <input checked="" type="checkbox"/> applicant and inventor <input type="checkbox"/> inventor only (If this check-box is marked, do not fill in below.) Applicant's registration No. with the Office
State (that is, country) of nationality: CN	State (that is, country) of residence: CN
This person is applicant for the purposes of: <input type="checkbox"/> all designated States <input type="checkbox"/> all designated States except the United States of America <input checked="" type="checkbox"/> the United States of America only <input type="checkbox"/> the States indicated in the Supplemental Box	
<input checked="" type="checkbox"/> Further applicants and/or (further) inventors are indicated on a continuation sheet.	
Box No. IV AGENT OR COMMON REPRESENTATIVE; OR ADDRESS FOR CORRESPONDENCE	
The person identified below is hereby/has been appointed to act on behalf of the applicant(s) before the competent International Authorities as: <input checked="" type="checkbox"/> agent <input type="checkbox"/> common representative	
Name and address: (Family name followed by given name; for a legal entity, full official designation. The address must include postal code and name of country.) China Patent Agent (H.K.) Ltd. 22/F, Great Eagle Centre 23 Harbour Road, Wanchai Hong Kong Special Administrative Region The People's Republic of China	Telephone No. (852)28284688 Facsimile No. (852)28271018 Teleprinter No. Agent's registration No. with the Office
<input type="checkbox"/> Address for correspondence: Mark this check-box where no agent or common representative is/has been appointed and the space above is used instead to indicate a special address to which correspondence should be sent.	

Box No. V DESIGNATIONS

The filing of this request constitutes under Rule 4.9(a), the designation of all Contracting States bound by the PCT on the international filing date, for the grant of every kind of protection available and, where applicable, for the grant of both regional and national patents.

However,

- ☐ DE Germany is not designated for any kind of national protection
- ☐ KR Republic of Korea is not designated for any kind of national protection
- ☐ RU Russian Federation is not designated for any kind of national protection

(The check-boxes above may be used to exclude (irrevocably) the designations concerned in order to avoid the ceasing of the effect, under the national law, of an earlier national application from which priority is claimed. See the Notes to Box No. V as to the consequences of such national law provisions in these and certain other States.)

Box No. VI PRIORITY CLAIM

The priority of the following earlier application(s) is hereby claimed:

Filing date of earlier application (day/month/year)	Number of earlier application	Where earlier application is:		
		national application: country or Member of WTO	regional application:* regional Office	international application: receiving Office
item (1)				
item (2)				
item (3)				

- ☐ Further priority claims are indicated in the Supplemental Box.

The receiving Office is requested to prepare and transmit to the International Bureau a certified copy of the earlier application(s) (only if the earlier application was filed with the Office which for the purposes of this international application is the receiving Office) identified above as:

- ☐ all items ☐ item (1) ☐ item (2) ☐ item (3) ☐ other, see Supplemental Box

* Where the earlier application is an ARIPO application, indicate at least one country party to the Paris Convention for the Protection of Industrial Property or one Member of the World Trade Organization for which that earlier application was filed (Rule 4.10(b)(ii)): . . .

Box No. VII INTERNATIONAL SEARCHING AUTHORITY

Choice of International Searching Authority (ISA) (if two or more International Searching Authorities are competent to carry out the international search, indicate the Authority chosen; the two-letter code may be used):

ISA / CN

Request to use results of earlier search; reference to that search (if an earlier search has been carried out by or requested from the International Searching Authority):

Date (day/month/year)

Number

Country (or regional Office)

Box No. VIII DECLARATIONS

The following declarations are contained in Boxes Nos. VIII (i) to (v) (mark the applicable check-boxes below and indicate in the right column the number of each type of declaration):

Number of
declarations

- | | | |
|---|--|---|
| <input type="checkbox"/> Box No. VIII (i) | Declaration as to the identity of the inventor | : |
| <input type="checkbox"/> Box No. VIII (ii) | Declaration as to the applicant's entitlement, as at the international filing date, to apply for and be granted a patent | : |
| <input type="checkbox"/> Box No. VIII (iii) | Declaration as to the applicant's entitlement, as at the international filing date, to claim the priority of the earlier application | : |
| <input type="checkbox"/> Box No. VIII (iv) | Declaration of inventorship (only for the purposes of the designation of the United States of America) | : |
| <input type="checkbox"/> Box No. VIII (v) | Declaration as to non-prejudicial disclosures or exceptions to lack of novelty | : |

Box No. IX CHECK LIST; LANGUAGE OF FILING

This international application contains:

(a) in paper form, the following number of sheets:

request (including declaration sheets) : 3
 description (excluding sequence listing and/or tables related thereto) : 13
 claims : 8
 abstract : 1
 drawings : 5

Sub-total number of sheets : 30

sequence listing :

tables related thereto :

(for both, actual number of sheets if filed in paper form, whether or not also filed in computer readable form; see (c) below)

Total number of sheets : 30

(b) ☐ only in computer readable form (Section 801(a)(i))(i) ☐ sequence listing(ii) ☐ tables related thereto(c) ☐ also in computer readable form (Section 801(a)(ii))(i) ☐ sequence listing(ii) ☐ tables related thereto

Type and number of carriers (diskette, CD-ROM, CD-R or other) on which are contained the

☐ sequence listing:☐ tables related thereto:

(additional copies to be indicated under items 9(ii) and/or 10(ii), in right column)

This international application is accompanied by the following item(s) (mark the applicable check-boxes below and indicate in right column the number of each item):

1. ☒ fee calculation sheet : 1
 2. ☒ original separate power of attorney : 1
 3. ☐ original general power of attorney :
 4. ☐ copy of general power of attorney; reference number, if any: :
 5. ☐ statement explaining lack of signature :
 6. ☐ priority document(s) identified in Box No. VI as item(s): :
 7. ☐ translation of international application into (language): :
 8. ☐ separate indications concerning deposited microorganism or other biological material :
 9. ☐ sequence listing in computer readable form (indicate type and number of carriers)
 (i) ☐ copy submitted for the purposes of international search under Rule 13ter only (and not as part of the international application) :
 (ii) ☐ (only where check-box (b)(i) or (c)(i) is marked in left column) additional copies including, where applicable, the copy for the purposes of international search under Rule 13ter :
 (iii) ☐ together with relevant statement as to the identity of the copy or copies with the sequence listing mentioned in left column :
 10. ☐ tables in computer readable form related to sequence listing (indicate type and number of carriers)
 (i) ☐ copy submitted for the purposes of international search under Section 802(b-quater) only (and not as part of the international application) :
 (ii) ☐ (only where check-box (b)(ii) or (c)(ii) is marked in left column) additional copies including, where applicable, the copy for the purposes of international search under Section 802(b-quater) :
 (iii) ☐ together with relevant statement as to the identity of the copy or copies with the tables mentioned in left column :
 11. ☐ other (specify): :

Figure of the drawings which should accompany the abstract:

Language of filing of the international application: EN

Box No. X SIGNATURE OF APPLICANT, AGENT OR COMMON REPRESENTATIVE

Next to each signature, indicate the name of the person signing and the capacity in which the person signs (if such capacity is not obvious from reading the request).



For receiving Office use only

1. Date of actual receipt of the purported international application: 23 · DEC 2005 (23 · 12 · 2005)	2. Drawings: <input type="checkbox"/> received: <input type="checkbox"/> not received:
3. Corrected date of actual receipt due to later but timely received papers or drawings completing the purported international application:	
4. Date of timely receipt of the required corrections under PCT Article 11(2):	
5. International Searching Authority (if two or more are competent): ISA /	6. <input type="checkbox"/> Transmittal of search copy delayed until search fee is paid

For International Bureau use only

Date of receipt of the record copy by the International Bureau:

This sheet is not part of and does not count as a sheet of the international application.

PCT

FEE CALCULATION SHEET

Annex to the Request

For receiving Office use only

PCT/CN 2005 / 0 0 2 3 0 5
International Application No.

23 · DEC 2005 (2 3 · 1 2 · 2 0 0 5)
Date stamp of the receiving Office

Applicant's or agent's
file reference

FPEL05150069

Applicant

INTEL CORPORATION etc.

CALCULATION OF PRESCRIBED FEES

1. TRANSMITTAL FEE

CNY500

T

CNY 500.-

2. SEARCH FEE

CNY1500

S

CNY 1500.-

International search to be carried out by

CN

(If two or more International Searching Authorities are competent to carry out the international search, indicate the name of the Authority which is chosen to carry out the international search.)

3. INTERNATIONAL FILING FEE

Where items (b) and/or (c) of Box No. IX apply, enter Sub-total number of sheets } 30
Where items (b) and (c) of Box No. IX do not apply, enter Total number of sheets }

i1 first 30 sheets

CHF1400

i1

CHF 1400.-

i2

number of sheets
in excess of 30

x

fee per sheet

=

i2

i3

additional component (only if sequence listing and/or tables related thereto are filed in computer readable form under Section 801(a)(i), or both in that form and on paper, under Section 801(a)(ii)):

400 x

fee per sheet

=

i3

Add amounts entered at i1, i2 and i3 and enter total at I

CHF1 400

I

CHF 1400.-

(Applicants from certain States are entitled to a reduction of 75% of the international filing fee. Where the applicant is (or all applicants are) so entitled, the total to be entered at I is 25% of the international filing fee.)

4. FEE FOR PRIORITY DOCUMENT (if applicable)

P

CNY 2000.-

5. TOTAL FEES PAYABLE

CNY2000CHF1400

TOTAL

CHF 1400.-

Add amounts entered at T, S, I and P, and enter total in the TOTAL box

MODE OF PAYMENT



authorization to charge
deposit account (see below)



postal money order



cash



coupons



cheque



bank draft



revenue stamps



other (specify):

AUTHORIZATION TO CHARGE (OR CREDIT) DEPOSIT ACCOUNT
(This mode of payment may not be available at all receiving Offices)



Authorization to charge the total fees indicated above.



(This check-box may be marked only if the conditions for deposit accounts of the receiving Office so permit) Authorization to charge any deficiency or credit any overpayment in the total fees indicated above.



Authorization to charge the fee for priority document.

Receiving Office: RO


Deposit Account No.

Date:

12/23/2005

Name:

Signature:



CHANGING A SCHEDULER IN A VIRTUAL MACHINE MONITOR

BACKGROUND

[0001] A virtual machine (VM) architecture logically partitions a physical machine, such that the underlying hardware of the machine is time-shared and appears
5 as one or more independently operation virtual machines. A computer platform in a virtual machine environment may comprise a virtual machine monitor (VMM) that may create a plurality of virtual machines and runs on the computer platform to facilitate for other software the abstraction of one or more virtual machines.

10[0002] The virtual machine monitor may comprise a scheduler to allocate time slots for each virtual machine to run and prioritize or balance the resource usage among the virtual machines. Usually, a scheduler may implement a specific scheduling mechanism that may fit specific situations, such as Borrowed Virtual Time (BVT) algorithm, Round Robin algorithm, etc.

15 BRIEF DESCRIPTION OF THE DRAWINGS

[0003] The invention described herein is illustrated by way of example and not by way of limitation in the accompanying figures. For simplicity and clarity of illustration, elements illustrated in the figures are not necessarily drawn to scale. For example, the dimensions of some elements may be exaggerated
20 relative to other elements for clarity. Further, where considered appropriate, reference labels have been repeated among the figures to indicate corresponding or analogous elements.

[0004] Fig. 1 shows an embodiment of a computer platform having a virtual machine monitor to change a scheduler.

[0005] Fig. 2 shows an embodiment of a scheduler manager in the virtual machine monitor of Fig. 1.

5[0006] Fig. 3 shows an embodiment of a method of changing a scheduler in the virtual machine monitor of Fig. 1.

[0007] Fig. 4 shows an embodiment of a method of handling a scheduling request by the scheduler changed in Fig. 3.

[0008] Fig. 5 shows an embodiment of a general computer platform having the
10 virtual machine monitor of Fig. 1.

DETAILED DESCRIPTION

[0009] The following description describes techniques for changing a scheduler in a virtual machine monitor. In the following description, numerous specific details such as logic implementations, pseudo-code, means to specify
15 operands, resource partitioning/sharing/duplication implementations, types and interrelationships of system components, and logic partitioning/integration choices are set forth in order to provide a more thorough understanding of the current invention. However, the invention may be practiced without such specific details. In other instances, control structures, gate level circuits and full
20 software instruction sequences have not been shown in detail in order not to obscure the invention. Those of ordinary skill in the art, with the included descriptions, will be able to implement appropriate functionality without undue experimentation.

[0010] References in the specification to "one embodiment", "an embodiment", "an example embodiment", etc., indicate that the embodiment described may include a particular feature, structure, or characteristic, but every embodiment may not necessarily include the particular feature, structure, or characteristic.

5 Moreover, such phrases are not necessarily referring to the same embodiment. Further, when a particular feature, structure, or characteristic is described in connection with an embodiment, it is submitted that it is within the knowledge of one skilled in the art to effect such feature, structure, or characteristic in connection with other embodiments whether or not explicitly
10 described.

[0011] Embodiments of the invention may be implemented in hardware, firmware, software, or any combination thereof. Embodiments of the invention may also be implemented as instructions stored on a machine-readable medium, that may be read and executed by one or more processors. A machine-readable
15 medium may include any mechanism for storing or transmitting information in a form readable by a machine (e.g., a computing device). For example, a machine-readable medium may include read only memory (ROM); random access memory (RAM); magnetic disk storage media; optical storage media; flash memory devices; electrical, optical, acoustical or other forms of
20 propagated signals (e.g., carrier waves, infrared signals, digital signals, etc.) and others.

[0012] An embodiment of a computer platform 10 having a virtual machine monitor to change a scheduler is shown in Fig. 1. The computer platform 10 may comprise an underlying hardware machine 11 having one or more processors
25 111, memory 112, console device 113, timer 114, and the like. The computer

platform 10 may further comprise a plurality of virtual machines and a virtual machine monitor. The plurality of virtual machines run their own operating systems and application software; such as a service virtual machine 13 running a service operating system 131 and a service application 132 and one or more guest virtual machines 14₁-14_N running their own guest operating systems 141₁-141_N and guest applications 142₁-142_N. The virtual machine monitor 12 may be responsible for processor(s)/memory virtualization/simulation, interrupt handling, virtual machine scheduling, etc. A non-exhaustive list of examples for the computer platform 10 may include mainframe computer, mini-computer, personal computer, portable computer, laptop computer and other devices for transceiving and processing data.

[0013] Various components in the computer platform 10 may invoke the virtual machine monitor 12 to perform a scheduling process, e.g., to determine a next owner of underlying switch devices for a next assigned time slot, or to determine a priority of a virtual machine, etc. Usually, the switch devices may be owned by a running virtual machine (e.g., a service virtual machine 13, or a guest virtual machine 14₁-14_N) in a time slot assigned by a scheduler in the virtual machine monitor and therefore have a focus on which virtual machine is currently active. Examples of the switch device may comprise the processor 111 and console device 113, wherein the console device 13 may comprise a frequently used I/O device, such as a keyboard, mouse, etc.

[0014] The service operating system 131 in the service virtual machine 13 may send a scheduling request for the scheduling process to the virtual machine monitor 12. For example, when the service operating system 131 monitors a running guest virtual machine 14₁-14_N, and finds a failure in the guest virtual

machine, e.g., a guest operating system fault, the service operating system 131 may send a scheduling request to the virtual machine monitor 12 to change the owner of the switch devices.

[0015] The guest operating system 141₁-141_N in the guest virtual machine 14₁-14_N may also send a scheduling request to the virtual machine monitor 12. For example, when the guest operating system 141₁-141_N is executing a device input/output operation and waiting for a response from the device, the guest operating system may send the scheduling request to yield the ownership of the switch devices so that other virtual machine may make a use of the switch devices.

[0016] The timer 114 in the underlying hardware, e.g., a programmable interval timer, may also send a scheduling request to the virtual machine monitor 12. For example, when the time slot assigned to the running virtual machine expires, the timer 114 may send a timer interrupt to the virtual machine monitor that may invoke the virtual machine monitor to perform the virtual machine scheduling in order to change the ownership of the switch devices.

[0017] The virtual machine monitor 12 may comprise a scheduler loader 120 and scheduler manager 121. The virtual machine monitor 12 may further comprise one or more schedulers; however, one of the one or more schedulers is active to handle the scheduling request. In the embodiment as depicted in Fig. 1, the virtual machine monitor 12 may comprise an old scheduler 122 and a new scheduler 123 that may respectively implement a specific scheduling mechanism such as Borrowed Virtual Time (BVT) algorithm, Round Robin algorithm, etc.

[0018] Many technologies may be applied as a working mode for the old scheduler 122 and the new scheduler 123. For example, the virtual machine monitor 12 may always hold the old scheduler 122 as a default scheduler. Before the new scheduler 123 is loaded in the virtual machine monitor 12, the old scheduler 122 may be active to handle the scheduling request. After the new scheduler 123 is loaded, the new scheduler may be active to handle the scheduling request in place of the old scheduler. However, if the new scheduler 123 is unloaded from the virtual machine monitor 12, the old scheduler 122 may be active again until the virtual machine monitor 12 is loaded with another new scheduler. For another example, in some circumstances, the scheduler manager 121 may switch back to the old scheduler 122 even though the new scheduler 123 exists in the virtual machine monitor 12. For still another example, the virtual machine monitor 12 may unload the old scheduler 122 before or after loading the new scheduler 123.

15[0019] The scheduler loader 120 may process a scheduler loading request from the service operating system 131 and load the new scheduler 123 in the virtual machine monitor 12 when the virtual machine monitor 12 or one or more of the virtual machines 13, 14₁-14_N is running. Software images of the new scheduler 123 may be available from various resources, such as a website, local disk, data center image server, etc. Example for an implementation of the scheduler loader 120 may comprises a hypercall handler that may process a hypercall for scheduler loading from the service operating system 131 and load the software images of the new scheduler 123 in the virtual machine monitor 12.

[0020] The scheduler manager 121 may be responsible for activating one of the old scheduler 122 and the new scheduler 123 to handle the scheduling request

when the virtual machine monitor 12 or one or more of the virtual machines 13, 14₁-14_N is running. The scheduler manager 121 may implement the scheduler activating in various ways. For example, the scheduler manager 121 may store a scheduler identifier to identify the scheduler that is active for the scheduling request. Before the new scheduler 123 is loaded in the virtual machine monitor 12, the scheduler manager 121 may store the old scheduler identifier in order to activate the old scheduler 122 to handle the scheduling request. After the new scheduler 123 is loaded, the scheduler manager 121 may replace the old scheduler identifier with the new scheduler identifier in order to activate the new scheduler to handle the scheduling request. However, when the new scheduler 123 is unloaded, the scheduler manager 121 may restore the old scheduler identifier to re-activate the old scheduler 122.

[0021] For another example, the scheduler manager may store a function pointer array pointing to a function array of the active scheduler, e.g., a function array of the old scheduler 122 or the new scheduler 123 which is active to handle the scheduling request. Fig. 2 depicts an embodiment for an implementation of the scheduler manager 121. As depicted, the scheduler manager 121 may comprise the function pointer array having a plurality of function pointers (e.g., pointers 0, 1, 2, 3, etc.). Each of the old scheduler 122 and new scheduler 123 may perform the scheduling process with a plurality of routine functions that adhere to a particular application programming interface (API), e.g., functions 0', 1', 2', etc. of the old scheduler 122 or functions 0'', 1'', 2'', etc. of the new scheduler 123. Each pointer in the function pointer array 121 may point to an active scheduler function, e.g., a function of the old scheduler 122 or the new scheduler 123. The scheduling requester, such as a virtual machine 13 or 14-

14_N, timer 114 or other devices that may trigger a scheduling process in the virtual machine monitor 12, may include a pointer to the function pointer array and may call the active scheduler function by dereferencing its pointer to the function pointer array and then calling the functions pointed by the array.

5[0022] In the embodiment depicted in Fig. 2, the scheduler manager 121 may activate the old scheduler 122 or the new scheduler 123 by updating the function pointer array to point to the active scheduler functions, or by updating the pointer of the scheduling requester to point to the function pointer array.

[0023] Referring back to Fig. 1, the scheduler manager 121 may be further
10 responsible for transporting information between the active scheduler (e.g., the old scheduler 122 or the new scheduler 123) and the scheduling requester. For example, in response to receiving the scheduling request from the scheduling requester, the scheduler manager 121 may dispatch the scheduling request to the active scheduler identified by the scheduler identifier or to the
15 active scheduler functions pointed by the function pointer array. The scheduler manager 121 may further send scheduling feedback information from the active scheduler to the scheduling requester. The scheduling feedback information may comprise a notification of the desired scheduling operation is performed correctly or not, a virtual machine priority information, and so on.

20[0024] However, other embodiments may implement other technologies for the structure of the computer platform 10. For example, the scheduler manager 121 may be omitted and the scheduling requester may issue a direct request to the virtual machine monitor requiring the scheduling process, wherein addresses associated to the active scheduler are dynamically patched into the

request during a scheduler model loading/unloading stage. With this means,
the request may be sent to the active scheduler directly.

[0025] Fig. 3 depicts an embodiment of a method of changing a scheduler in the
virtual machine monitor 12 as shown in Fig. 1. In block 301, a user or other
5 suitable party may decide to change an old scheduler (e.g., the old scheduler
122) currently active to handle a scheduling request in the virtual machine
monitor with a new scheduler (e.g., the scheduler 123) when the virtual
machine monitor or one or more of the virtual machines is running. For
example, a user may determine that the old scheduler is a poor fit for the
10 current virtualization environment or that another schedule is a better fit for the
current virtualization environment.

[0026] In block 302, the decision made in block 301 may invoke an application
running over a service operating system in a service virtual machine (e.g.,
service operating system 131) with specified parameters and the application
15 may pass the parameters and other information to the service operating
system through a virtual machine control request, that may trigger the service
operating system to issue a scheduler loading request into the virtual machine
monitor to load the new scheduler. The specific parameters may comprise
information on which scheduler is to be loaded, where the scheduler image is;
20 what kind of loading policy the virtual machine monitor may apply to load the
new scheduler, etc.

[0027] In block 303, a scheduler loader (e.g., the loader 120) or other suitable
component may handle the loading request and cease all of switch devices
owned by a running virtual machine, wherein the switch devices may comprise
25 processor(s) and console devices. Various methods may be applied to perform

the ceasing process. For processor ceasing, a cease sign may be input in all of virtual machine resume paths between all of virtual processor(s) in the virtual machine monitor and the running virtual machine. For console device ceasing, the virtual machine monitor may flush all of outstanding traffic by the help of console device models before the virtual machine monitor actually stop them. After the switch devices are ceased, the switch devices may reach a stable, consistent or predictable state so that the virtual machine monitor may retain their state for a next scheduling process.

[0028] In block 304, the scheduler loader or other suitable component may decide whether to unload the old scheduler from the virtual machine monitor before loading the new scheduler. In different circumstances, the scheduler loader may make different decision. For example, the scheduler loader may decide to reserve the old scheduler for a future use. However, if there is no free space for the new scheduler, the scheduler loader may decide to unload the old scheduler. In response to deciding to reserve the old scheduler, the scheduler loader or other suitable component may load the new scheduler in the virtual machine monitor with use of the parameters from the service operating system in block 306. However, in response to deciding to unload the old scheduler, the scheduler loader or other suitable component may unload the old scheduler from the virtual machine monitor in block 305 and then load the new scheduler in the virtual machine monitor in block 306. Other embodiments may implement other technologies for the scheduler unloading. For example, the scheduler loader may make a decision on whether to unload the old scheduler after the new scheduler is loaded.

[0029] In block 307, the new scheduler may be activated to handle the scheduling request in place of the old scheduler. For example, the scheduler manager 121 may activate the new scheduler by means of storing an identifier to identify the new scheduler or a function pointer array pointing to the new scheduler functions. For another example, the new scheduler may be activated by dynamically patching addresses associated with the new scheduler in the scheduler request so that a scheduling requester may directly call the new scheduler to perform the scheduling process.

[0030] Fig. 4 depicts an embodiment of a method of handling a scheduling request by the new scheduler changed in Fig. 3. In block 401, the virtual machine monitor may receive the scheduling request from the scheduling requester requiring a scheduling process. In block 402, the scheduling request may be transferred to the new scheduler, for example, through a scheduler manager (e.g., the scheduler manager 121) that may store an identifier to identify the new scheduler or a function pointer array pointing to a function array of the new scheduler.

[0031] In block 403, the new scheduler may handle the scheduling request, for example, determine a next owner of the switch devices or calculate a scheduling priority for a specific virtual machine. In block 404, the new scheduler may return scheduling feedback information to the scheduling requester. The scheduling feedback information may comprise a notification of the desired scheduling process is performed correctly or not, a virtual machine priority information, and so on.

[0032] Fig. 5 depicts an embodiment of a general computer platform having the virtual machine monitor as shown in Fig. 1. The computing platform may

comprise one or more processors 50, memory 51, chipset 52, I/O device 53, BIOS firmware 54 and the like. The one or more processors 50 are communicatively coupled to various components (e.g., the memory 51) via one or more buses such as a processor bus. The processors 50 may be implemented as an integrated circuit (IC) with one or more processing cores that may execute codes under a suitable architecture, for example, including Intel® Xeon™, Intel® Pentium™, Intel® Itanium™ architectures, available from Intel Corporation of Santa Clara, California.

[0033] In an embodiment, the memory 51 may store codes to be executed by the processor 50. A non-exhaustive list of examples for the memory 51 may comprise one or a combination of the following semiconductor devices, such as synchronous dynamic random access memory (SDRAM) devices, RAMBUS dynamic random access memory (RDRAM) devices, double data rate (DDR) memory devices, static random access memory (SRAM), flash memory devices, and the like.

[0034] In an embodiment, the chipset 52 may provide one or more communicative path among the processor 50, memory 51 and various components, such as the I/O device 53 and BIOS firmware 54. The chipset 52 may comprise a memory controller hub 520, an input/output controller hub 521 and a firmware hub 522.

[0035] In an embodiment, the memory controller hub 520 may provide a communication link to the processor bus that may connect with the processor 50 and to a suitable device such as the memory 51. The memory controller hub 520 may couple with the I/O controller hub 521 that may provide an interface to the I/O devices 53 for the computing platform such as a keyboard

and a mouse. A non-exhaustive list of examples for the I/O devices 13 may comprise a keyboard, mouse, network card, a storage device, a camera, a blue-tooth, an antenna, and the like.

[0036] In an embodiment, the memory controller hub 520 may communicatively couple with a firmware hub 522 via the input/output controller hub 521. The firmware hub 522 may couple with the BIOS firmware 54 that may store routines that the computing platform executes during system startup in order to initialize the processors 50, chipset 52, and other components of the computing platform. Moreover, the BIOS firmware 54 may comprise routines or drivers that the computing device 1 may execute to communicate with one or more components of the compute platform.

[0037] The computer platform as depict in Fig. 5 may perform as the computer platform 10 as depicted in Fig. 1. The memory 51 may store software images as a virtual machine monitor including a scheduler loader, one or more scheduler, and/or a scheduler manager. The memory 51 may further store service software including service operating system and service applications, and guest software including guest operating system and guest applications.

[0038] While certain features of the invention have been described with reference to example embodiments, the description is not intended to be construed in a limiting sense. Various modifications of the example embodiments, as well as other embodiments of the invention, which are apparent to persons skilled in the art to which the invention pertains are deemed to lie within the spirit and scope of the invention.

What is claimed is:

1. A method for changing a first scheduler in a virtual machine monitor,
comprising:

loading a second scheduler in the virtual machine monitor when the virtual
5 machine monitor is running; and

activating the loaded second scheduler to handle a scheduling request for
a scheduling process in place of the first scheduler when the virtual machine
monitor is running.

10 2. The method of claim 1, wherein the loading further comprises:
ceasing device resources owned by a running virtual machine in response
to receiving a scheduler changing request to change the first scheduler; and
loading the second scheduler in the virtual machine monitor based upon a
scheduler parameter of the scheduler changing request.

15 3. The method of claim 1, wherein the loading further comprises:
unloading the first scheduler from the virtual machine monitor before
loading the second scheduler.

20 4. The method of claim 1, wherein the activating further comprises:
replacing a first scheduler identifier with a second scheduler identifier to
route between the second scheduler and a requester that generated the
scheduling request, when the virtual machine monitor is running.

5. The method of claim 1, wherein the activating further comprises:

replacing a first function pointer array pointing to a first function array of the first scheduler with a second function pointer array pointing to a second function array of the second scheduler to route between the second scheduler and a requester that generated the request, when the virtual machine monitor is running.

5

6. The method of claim 1, wherein the activating further comprises:

dynamically patching an address associated with the second scheduler into the scheduling request when the virtual machine monitor is running.

10

7. The method of claim 1, further comprising:

unloading the second scheduler from the virtual machine monitor when the virtual machine monitor is running; and

re-activating the first scheduler to handle a scheduling request after the second scheduler has been unloaded.

15

8. A virtual machine monitor for changing a first scheduler, comprising:

a loading logic to load a second scheduler in the virtual machine monitor when the virtual machine monitor is running; and

an activating logic to activate the loaded second scheduler to handle a scheduling request for a scheduling process in place of the first scheduler when the virtual machine monitor is running.

20

9. The virtual machine monitor of claim 8, wherein the loading logic is further to:

cease device resources owned by a running virtual machine in response to receiving a scheduler changing request to change the first scheduler; and

load the second scheduler in the virtual machine monitor based upon a scheduler parameter of the scheduler changing request.

5

10. The virtual machine monitor of claim 8, wherein the loading logic is further to:

unload the first scheduler from the virtual machine monitor before loading the second scheduler.

10

11. The virtual machine monitor of claim 8, wherein the activating logic is further to:

replace a first scheduler identifier with a second scheduler identifier;

route between the second scheduler as identified by the second scheduler identifier and a requester that generated the scheduling request, when the virtual machine monitor is running.

15

12. The virtual machine monitor of claim 8, wherein the activating logic is further to:

replace a first function pointer array pointing to a first function array of the first scheduler with a second function pointer array pointing to a second function array of the second scheduler;

20

route between the second function array pointed by the second function pointer array and a requester that generated the scheduling request, when the

virtual machine monitor is running.

25

13. The virtual machine monitor of claim 8, wherein the activating logic is further to:

5 dynamically patch an address associated with the second scheduler into a scheduling request when the virtual machine monitor is running.

14. The virtual machine monitor of claim 8, wherein the loading logic is further to unload the second scheduler from the virtual machine monitor when the virtual machine monitor is running; and the activating logic is further to re-activate
10 the first scheduler to handle a scheduling request after the second scheduler has been unloaded.

15. A system, comprising:
a requester to generate a scheduling request for a scheduling process;
15 a virtual machine monitor, comprising:
 a loading logic to load a second scheduler in the virtual machine monitor when the virtual machine monitor is running; and
 an activating logic to activate the loaded second scheduler to handle the scheduling request in place of a first scheduler when the virtual
20 machine monitor is running.

16. The system of claim 15, wherein the requester further comprises at least one of a timer, a service virtual machine and a guest virtual machine.

17. The system of claim 15, wherein the requester is further to generate a scheduler changing request to changing the first scheduler.

18. The system of claim 15, wherein the loading logic is further to:
5 cease device resource owned by a running virtual machine in response to receiving a scheduler changing request to change the first scheduler; and
 load the second scheduler in the virtual machine monitor based upon a scheduler parameter of the scheduler changing request.

19. The system of claim 15, wherein the loading logic is further to:
10 unload the first scheduler from the virtual machine monitor before the second scheduler is loaded.

20. The system of claim 15, wherein the activating logic is further to:
15 replace a first scheduler identifier with a second scheduler identifier;
 route between the second scheduler as identified by the second scheduler identifier and the requester, when the virtual machine monitor is running.

21. The system of claim 15, wherein the activating logic is further to:
20 replace a first function pointer array pointing to a first function array of the first scheduler with a second function pointer array pointing to a second function array of the second scheduler;
 route between the second function array pointed by the second function pointer array and the requester, when the virtual machine monitor is running.

25

21)

22. The system of claim 15, wherein the activating logic is further to:

dynamically patch an address associated to the second scheduler to the scheduling request when the virtual machine monitor is running.

5 23. The system of claim 15, wherein the loading logic is further to unload the second scheduler from the virtual machine monitor when the virtual machine monitor is running; and the activating logic is further to re-activate the first scheduler to handle the scheduling request after the second scheduler has been unloaded.

10

24. A machine readable medium comprising a plurality of instructions that in response to being executed result in an apparatus:

loading a second scheduler in a virtual machine monitor when the virtual machine monitor is running; and

15 activating the loaded second scheduler to handle a scheduling request for a scheduling process in place of a first scheduler, when the virtual machine monitor is running.

20 25. The machine readable medium of claim 24, wherein the plurality of instructions that result in the apparatus loading the second scheduler, further result in the apparatus:

ceasing device resources owned by a running virtual machine in response to receiving a scheduler changing request to change the first scheduler; and loading the second scheduler in the virtual machine monitor based upon a scheduler parameter of the scheduler changing request.

5

26. The machine readable medium of claim 24, wherein the plurality of instructions further result in the apparatus:

unloading the first scheduler from the virtual machine monitor before the second scheduler is loaded.

10

27. The machine readable medium of claim 24, wherein the plurality of instructions that result in the apparatus activating the second scheduler, further result in the apparatus:

replacing a first scheduler identifier with a second scheduler identifier to route between the second scheduler and a requester that generated the scheduling request, when the virtual machine monitor is running.

15

28. The machine readable medium of claim 24, wherein the plurality of instructions that result in the apparatus activating the second scheduler, further result in the apparatus:

20

replacing a first function pointer array pointing to a first function array of the first scheduler with a second function pointer array pointing to a second function array of the second scheduler to route between the second scheduler and a requester that generated the scheduling request, when the virtual machine
5 monitor is running.

29. The machine readable medium of claim 24, wherein the plurality of instructions that result in the apparatus activating the second scheduler, further result in the apparatus:
10 dynamically patching an address associated with the second scheduler to the scheduling request when the virtual machine monitor is running.

30. The machine readable medium of claim 24 wherein the plurality of instructions further result in the apparatus:
15 unloading the second scheduler from the virtual machine monitor when the virtual machine monitor is running; and
re-activating the first scheduler to handle the scheduling request after the second scheduler has been unloaded.

ABSTRACT

Machine-readable media, methods, and apparatus are described to change a first scheduler in the virtual machine monitor. In some embodiments, a second scheduler is loaded in a virtual machine monitor when the virtual machine monitor is running; and then is activated to handle a scheduling request for a scheduling process in place of the first scheduler, when the virtual machine monitor is running.

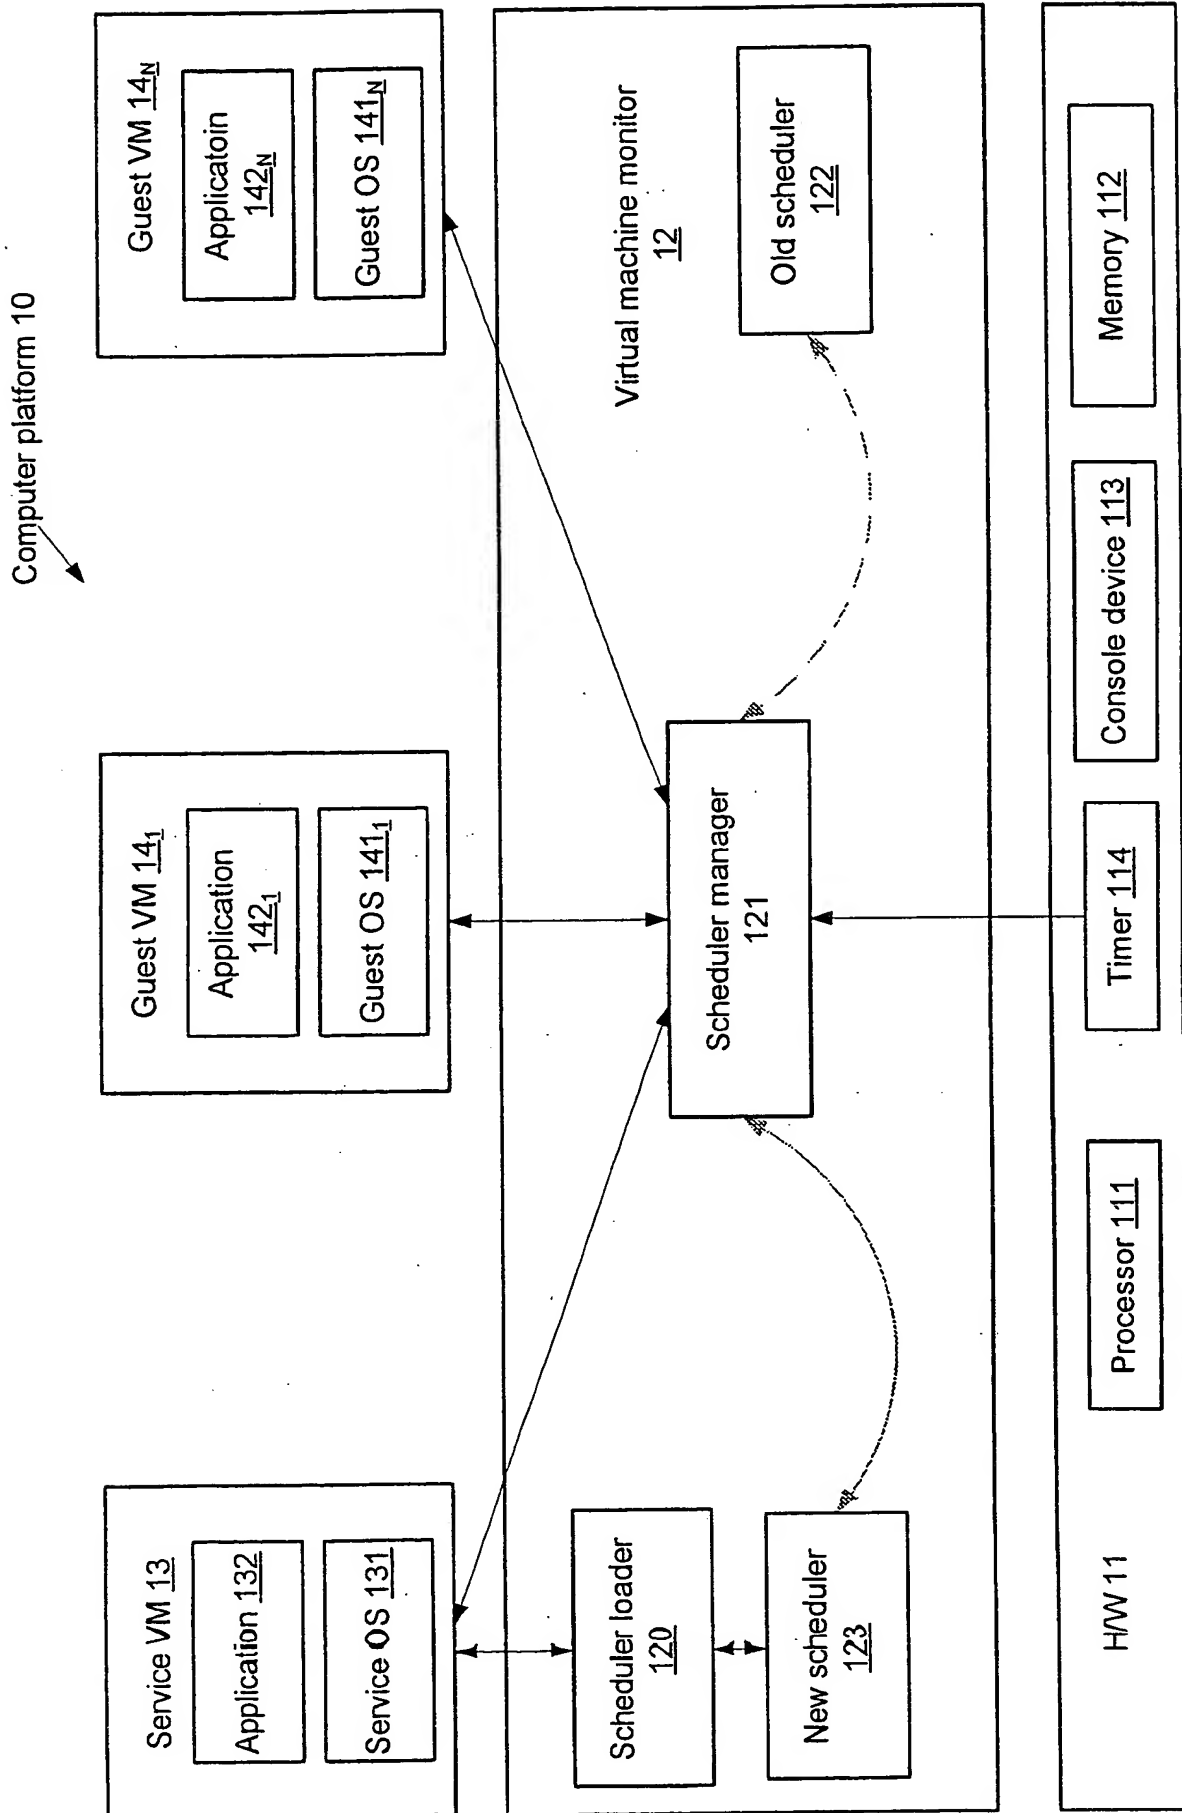


FIG. 1

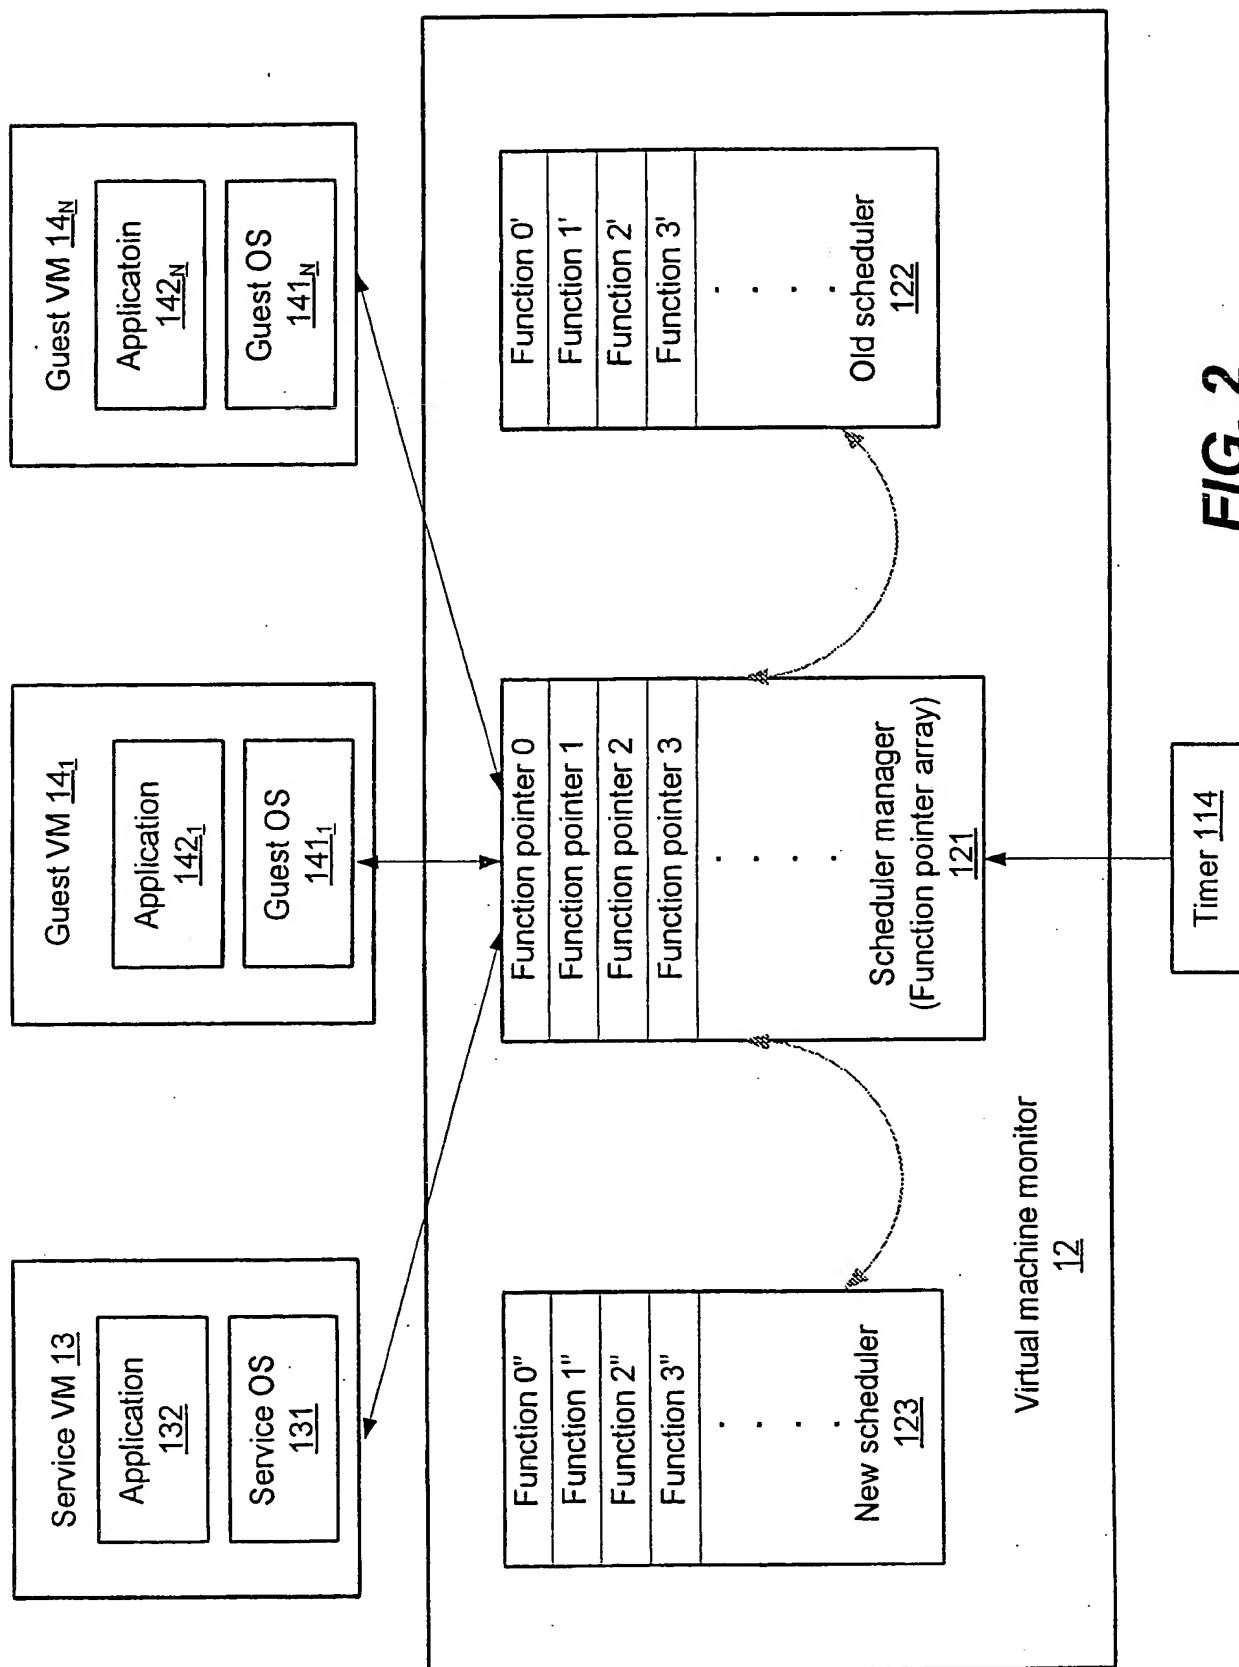
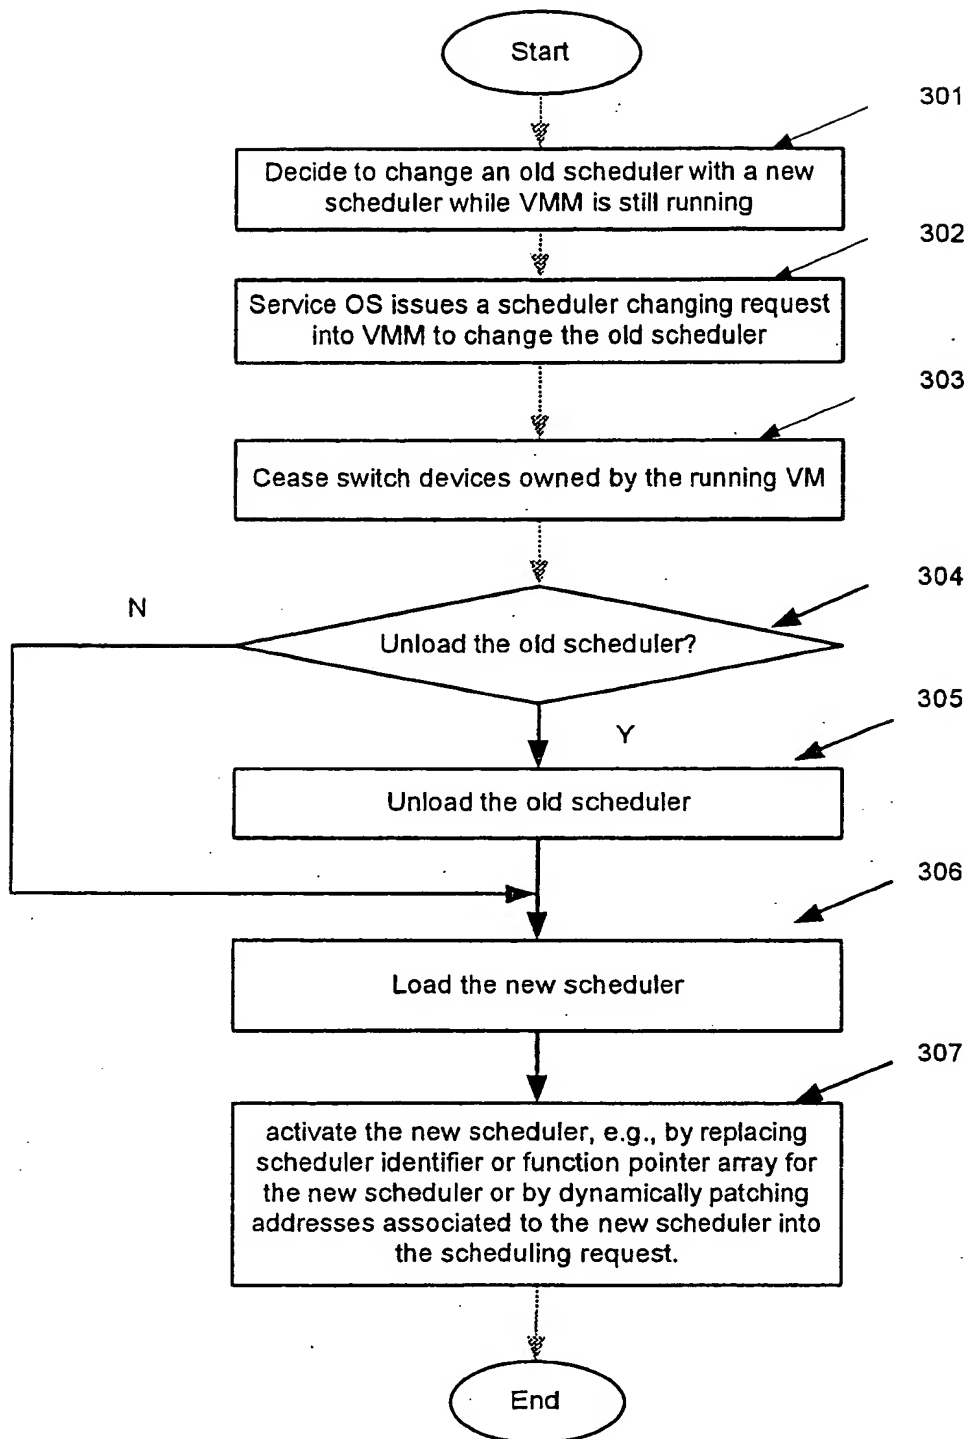
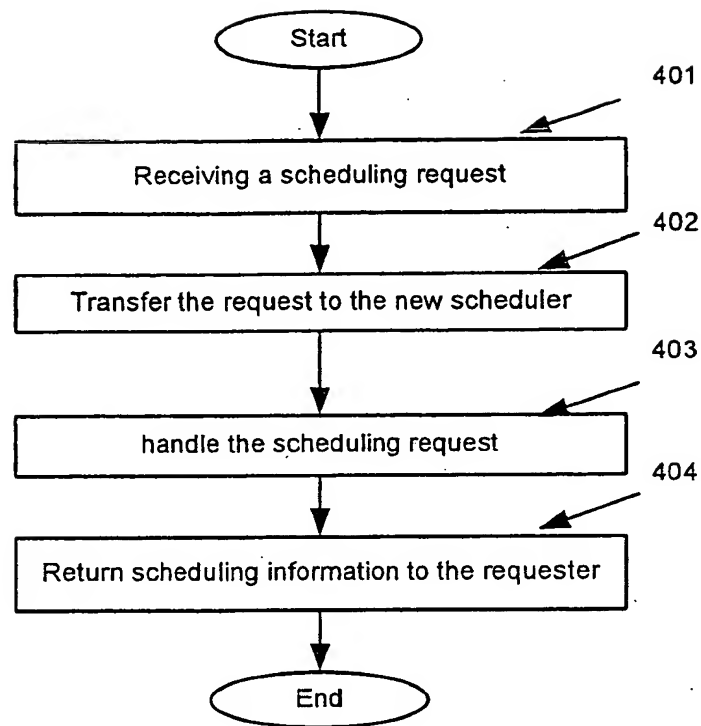
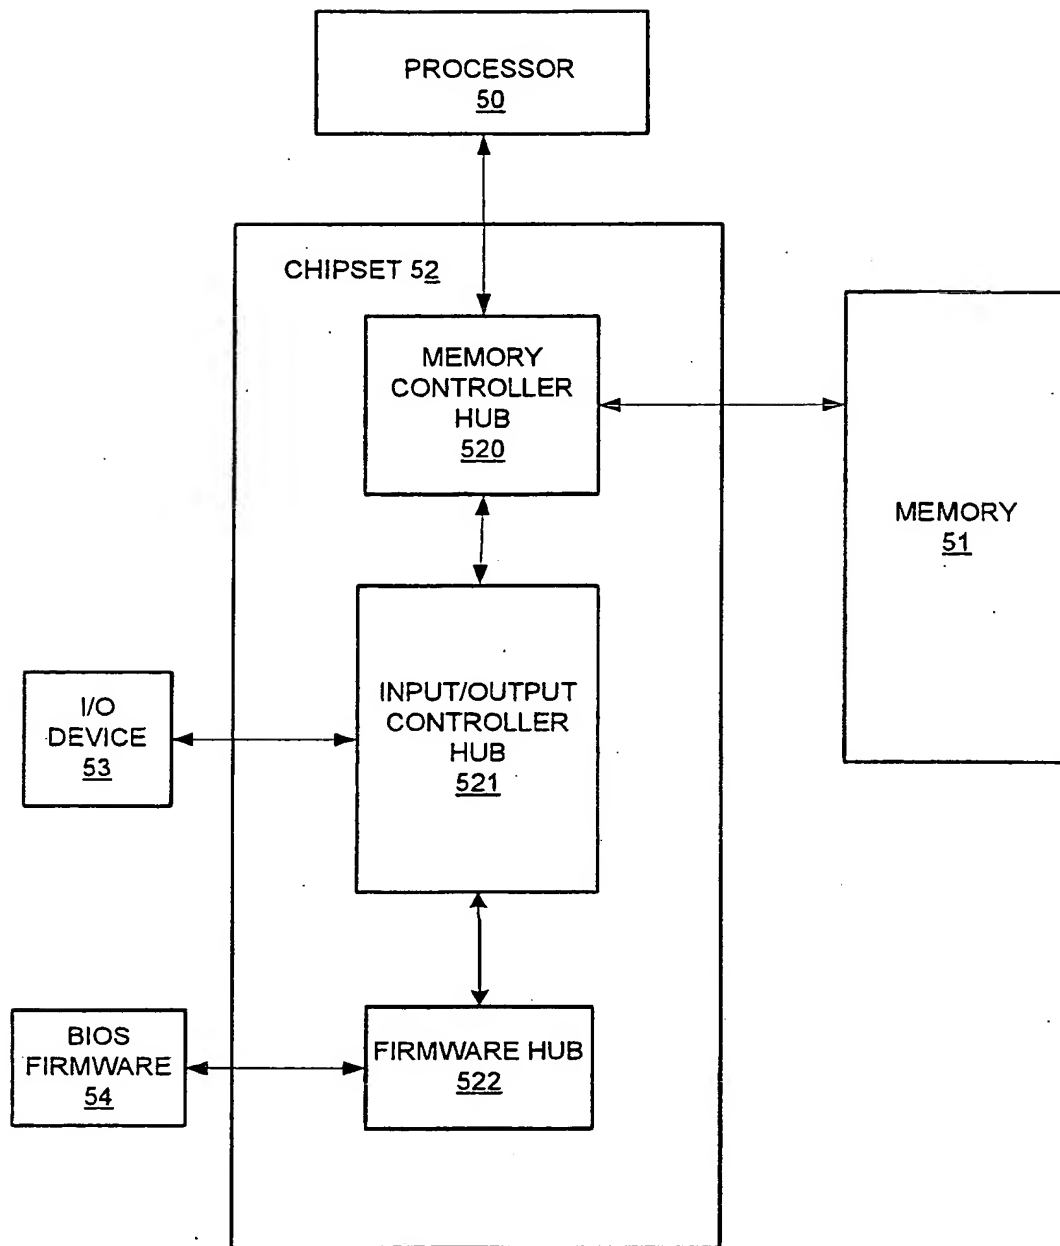


FIG. 2

**FIG. 3**

**FIG. 4**

5/5

**FIG. 5**